

REMARKS

In response to the Official Action dated March 26, 2003, Applicants respectfully request reconsideration and withdrawal of the rejections.

General Overview

The Applicants' invention provides an apparatus and method that are suitable in particular for streaming data, especially live real-time data. A primary objective of the Applicants' invention is to improve the quality of streaming data without requiring very large bandwidth. This is achieved, in summary, by providing a network of gateways connecting individual clients to the servers that are the sources of streaming. If a gateway is already providing a given stream to one client, then upon receipt of a request for that stream from a second client, the gateway is able to supply the same stream to the second client by "copying" the stream using a unicast-multicast-unicast scheme within the gateway.

If the gateway receives a request for a stream that it is not already supplying to another client, then the gateway may look for a neighboring gateway that is already handling this stream and which may be used as the source of the stream, rather than having to request the stream from the server.

The Applicants' invention is advantageous in that by "copying" streams within a gateway, and by looking to other gateways as possible sources for the stream in place of the originating server, the demand on the server bandwidth, and network bandwidth, is reduced. It should be noted here that the stream is not literally "copied", but is made

available to more than one client by changing the packet address-type from unicast to multicast within the gateway, and by then changing the address-type back to unicast if (as would be common) the clients are not in a multicast-enabled environment.

The rejections will be discussed hereinafter, with reference to the item numbers appearing in the Action.

Claim Rejections - 35 USC 112

1-2. Claims 13 and 25 have been amended as suggested by the examiner.

Claim Rejections - 35 USC 102

5. Claim 27 was rejected as being anticipated by the Shur et al. patent. To clarify the distinctions from this reference, the claim has been amended to recite that the gateways are capable of automatically changing the address of a data packet in four different ways. It is respectfully submitted that the Shur patent does not anticipate this subject matter.

6. Claim 28 was rejected on the grounds that it was considered to be anticipated by the Yates et al. patent. Yates discloses an arrangement requiring a network of cache servers. It should be noted, however, that the Applicants' invention is not restricted to cached content. In particular the Applicants' invention relates especially to live content that the gateway is currently serving to some clients. In Applicants' invention, if the gateway

does not have the requested content, it may be obtained from any other neighboring gateway which can provide the stream with the highest quality (which may change dynamically). In Yates, however, if a gateway does not have the content, then the search for other gateways that might have the content is restricted to the next router along the path to the home server (column 7, line 42-45) and there is no guarantee that this router could necessarily provide the content with a high quality.

Claim Rejections – 35 USC 103

8-9. Claims 1, 2, 14 and 15 were rejected as being unpatentable over the Burns et al patent in view of the Haggerty et al. patent. Burns et al. teaches how an ISP cache server can stream the cached audio/video data to the user to improve performance. The contents are downloaded from the original content server to the ISP cache server during the off-peak hours in advance. This is similar to mirror sites in which contents are mirrored to the content servers. It should be noted that the approach of Burns requires a prediction to be made in advance of the content required, which is unlikely to be very accurate. This is in contrast to the Applicants' invention which is not a mirror or cache server approach, but which provides a network in which a gateway can search for the content from the nearest source dynamically as and when required.

Haggerty teaches a means for supplying a second and subsequent client with a data stream already being supplied to a first client. In Haggerty, this is discussed within the context of IP multicast, a network-layer standard defined by IETF more than 10 years ago.

The Applicants' invention does not use IP multicast to achieve the same purpose of supplying a second and subsequent client with a data stream already being supplied to a first client. The gateways of the Applicants' invention achieve efficient one-to-many (similar but not the same as IP multicast) without relying on IP multicast capability which may be lacking in some network environments such as conventional Internet Service Providers, and indeed the Internet generally.

The Haggerty patent is specifically concerned with the ability to provide multicast traffic in a switch-based network. See column 7, lines 5-20. There is no apparent reason why a person of ordinary skill in the art would be motivated to apply its teachings to the system of the Burns patent. Specifically, there is no teaching in either reference that suggests the use of multicast transmission for the delivery of streaming data from a content provider. Only Applicants' disclosure provides a teaching of such a concept.

Accordingly, it is respectfully submitted that it would not be obvious to combine these references, absent knowledge of the present invention. Reconsideration and withdrawal of the rejection is therefore requested.

10. Claims 2 and 15 are dependent claims from claims 1 and 14 respectively. Since claims 1 and 14 should be considered allowable, it follows that claims 2 and 15 are likewise allowable. The further rejection of these claims based on Burns is however traversed for the following reasons.

Again Burns teaches how a local cache server can serve the user request if it has a cached copy of the content stored locally. The main difference between Burns and claims 2 and 15 is that Applicants' invention as defined in these claims deals with non-cacheable contents in which streaming contents cannot be cached in advance to the local servers while Burns proposes to pre-cache popular contents to local cache servers.

11-13. Claims 3-5 and 16-18 were rejected on the basis of the Burns and Haggerty patents, in further view of the Shur patent. Shur's MUS (Multicast-Unicast Server) aims to allow unicast clients to connect to multicast sessions through the MUS. On one side of the MUS, the MUS joins a multicast session to receive while on the other side, the MUS uses individual unicast connections to serve each client. The purpose of Applicants' invention is different from MUS. Applicants' invention aims at achieving multicast at the application level without using IP multicast in either side of the network. In the Applicants' invention a gateway receives from one side a unicast session and duplicates the received information, which would be sent out to each individual client through a unicast session. This achieves multicasting at the application level without using IP multicast external to the gateway at either side. It is important to note that in the Applicants invention the duplication of received information is done internally within the gateway by using IP multicast only within the gateway. In other words, unicast is used external to the gateway for both receiving and forwarding to each client. IP multicast is used only internally within the gateway for the purpose of efficient duplication of received information.

14. The use of TTL as defined in the IETF standard and explained in many textbooks on computer networks is that an initial TTL value is attached to the header of a packet to limit how many routers or hops the packet can go through. When the packet is received at a router, its TTL value would be decremented by one before being forwarded to the next hop router. When the TTL value reaches the value of zero, the packet would be dropped by the router without being forwarded to the next hop. In the Applicants' invention IP multicast is used internally within the gateway for efficient duplication of received information, i.e., the received information will be virtually sent out using some IP multicast address that exists only within the gateway. Each client wishing to receive the duplicated information would join the multicast session by receiving from the multicast address. This achieves the purpose of packet duplication within the gateway by using IP multicast internally within the gateway. Since this packet duplication exists internally within the gateway in the Applicants' invention, the TTL value of the multicast packet must be set to zero. It should be noted that setting the TTL to zero is a counter-intuitive process since normally it would be regarded as a pointless thing to do. It is done in the present invention simply since the multicast takes place only within the gateway and it is necessary to prevent packets from being spread widely beyond the gateway.

15. Shur's MUS is to explain how interworking between unicast and multicast is achieved through the use of MUS at the boundary between the multicast domain where the multicast session exists and the unicast domains through which the clients are accessing the

Internet. Since IP multicast is used only internally within the gateway in the Applicants' invention for the purpose of efficient duplication of received information, unicast is normally used outside the gateway (unless the client is in a multicast domain) and thus a multicast to unicast conversion is needed after the duplication of packets. In other words, the Applicants' invention achieves application level multicast by using unicast sessions at the incoming and outgoing sides. IP multicast is used for the purpose of efficient duplication only. On the other hand, MUS enables network-level multicast even when the client is connected to unicast capable network domains. The Applicants' invention includes the possibility that the gateway may be provided within a unicast-only network, because multicast is only used within the gateway and at the application level.

16-18. Claims 6 and 19 were rejected on the basis of the Burns, Haggerty and Shur patents, in further view of the Luczycki et al. patent. Luczycki et al.'s network traffic exchange facility permits independent networks to exchange IP multicast data streams, which implies that the source's multicast address type is changed to another multicast address type. The network exchange facility operates in the IP layer to allow the exchange of IP multicast packets across different network domains through the change of IP multicast address at the network exchange facility.

The Applicants' invention on the other hand is different from the network exchange facility in that Applicants' invention operates at the application layer. Applicants' invention performs multicast based on the URL of content, not IP multicast address. Multicasting is

achieved at the application level among Applicants' invention over networks with or without IP multicast capability.

Applicants' invention can perform application-level multicast through any input-output combination (i.e., unicast → unicast, unicast → multicast, multicast → unicast, and multicast → multicast) among the gateways participating in the application-level multicast.

19-21. Claims 7, 8, 20 and 21 were rejected over the Yates patent in view of the Haggerty patent. As stated earlier, Yates restricts the cache servers to be found only along the path back to the home server. In addition, Yates requires the cache server to be able to insert a packet filter into the router associated with it in order to intercept requests for content.

Applicants' invention differs from Yates in at least two important aspects. First, the search of the required content is not restricted only along the path to the home server. Because the network conditions will be continually changing, there is no guarantee that going back along the path to the home server will provide the highest quality streaming. By providing this flexibility, Applicants' invention allows the gateway to search for the best path that could in all probability be different from the default path back to the home server, and indeed the "best path" may change dynamically and Applicants' invention allows the choice of best path to be continually changed. Secondly, the Applicants' invention does not require the insertion of any agent software (such as the packet filter) into the router

associated with it. It is not practical to insert software inside third-party routers to intercept user requests (especially in large-scale deployments and networks).

As explained above, Haggerty teaches a means for supplying a second and subsequent client with a data stream already being supplied to a first client. This is discussed within the context of IP multicast, a network-layer standard defined by IETF more than 10 years ago. The Applicants' invention uses gateways that do not use IP multicast to achieve the same purpose of supplying a second and subsequent client with a data stream already being supplied to a first client. The gateways in the Applicants' invention achieve efficient one-to-many (similar to IP multicast) without relying on an IP multicast capability which may be lacking in some network environments.

As with the rejection of claims 1 and 14, there are no teachings in either the Yates patent or the Haggerty patent that would lead a person of ordinary skill in the art to combine them. First, as is apparent from its cover figure, the Yates patent discloses a router-based network, and therefore would have no need for the specific teachings of Haggerty relating to switched networks. Second, and perhaps more importantly, it is not apparent why one would be inclined to use multicasting in the system of Yates. Only Applicants have provided any motivation for a combination of these two patents.

22. The neighboring gateways in Yates' include only those along the path back to the home server while in Applicants' invention the neighboring gateways can be along any paths to allow more flexibility and more choices in path selection. In addition, the

gateways in the Applicants' invention can handle streaming content that are not cacheable while Yates' cache servers deal only with cached contents.

23. Yates teaches only a method for passing the request up to the next router along the path to the home server if a cache copy is not encountered at the particular cache server. Applicants' invention differs from the router/cache server combination in Yates. First, Applicants' invention does not need to interwork with routers for request interception. Secondly, in the Applicants' invention a gateway searches for the required non-cacheable content from one of the neighboring gateways, and in the Applicants' invention this is not necessarily along the path back to the home server.

24. Claims 9, 10, 22 and 23 were rejected on the basis of the Yates and Haggerty patents in further view of the Lin et al. patent. Lin proposes a totally unrelated scheme ("when client device begins to process data segment one, it also sends a request signal to the appropriate upstream caching server requesting the data segment two. The appropriate upstream caching server was determined by server 100 during the initialization phase. Client device 120 sends a request signal to caching server 114 (level N), requesting the data segment two. Concurrently, caching server 114 (level N) sends a request signal to its upstream caching server 112 (level N-1) requesting the data segment three, ..."). The request signal is to inform the upstream caching server to send a data segment in anticipation of its use by a downstream caching server or client. The request signals are

sent only by a downstream caching server to the upstream caching server to request a particular data segment "n".

In the Applicants' invention a gateway reports to each neighboring gateway (not only the upstream caching server) regarding a new stream content being served by the current gateway. This is for the purpose of letting other gateways know about the content currently being served by each gateway so that other gateways may get the required contents from its neighbors having the content. It is not for the purpose of letting other gateways become aware that it is busy sending data streams to the client.

Applicants' invention has the capability to select between two or more possible gateways as the source of a data stream requested by a client. This is different from the teaching of Lin (column 7 line 66 – column 8 line 21 as described above involving caching server (level N), caching server (level N-1), etc.). In Lin, there is no choice of caching server. It is always the case that caching server (level N) send request signals to caching server (level N-1), etc.

25. As is explained above, in Lin the caching servers along a path are chosen during the initialization phase. After that, the caching servers send requests only to the upstream caching server for the purpose of requesting a particular data segment in advance. This is not for the purpose of informing all the neighboring gateways that the current gateway is providing the data stream so that they can get the data stream from the current gateway if they want the data stream in the future.

26. In Lin, the choice of neighboring caching server is always the upstream caching server along the pre-selected path (i.e., there is no choice in Lin's invention). In claim 23, the gateway can select from two or more choices.

27-28. Claims 11-13, 24 and 25 are rejected on the basis of the Yates, Haggerty and Lin references, in further view of the O'Neil et al. patent. O'Neil teaches how to load balance among a cluster of servers. This is very different from the Applicants' invention. O'Neil teaches a system for implementing peer-to-peer load balancing among plural network servers (column 4, lines 63-65). In the Applicants' invention the gateways are not a cluster of network servers. They are a set of gateways (servers) deployed across the Internet at distant locations. The latency between two gateways can vary a lot while the latency between two servers in a cluster of network servers can be either a very small value or infinity (when one of the servers is not responding). In addition to the difference in latency, Applicants' invention also monitors the quality of the streams supplied by possible source gateways. The quality of the streams in a cluster of network servers directly depends on the loading of the possible network server while in the case of gateways deployed over the Internet, the quality of the streams depends on both the loading of the possible source gateway and the bandwidth availability between the possible source gateway and the current gateway. It is therefore important to consider the quality of streams, not just the loading of the possible source gateway.

The rejection comments that a function of the Applicants' invention is to make the system more time efficient. In fact, the purpose of the Applicants' invention is not time efficiency but for the purpose of ensuring good streaming quality.

29. The criteria of prioritization among the quality of stream, the source gateway loading, and latency are important for streaming real-time data as in the Applicants' invention case and not as important in the case of WWW server load balancing. It is respectfully submitted that the rejection of claim 13 does not find any support in the applied references, and must therefore be based on Applicants' own teaching.

33. Claim 26 was rejected on the basis of the Shur patent in view of the Luczycki patent. For the reasons presented previously in connection with claims 6 and 19, it is respectfully submitted that the subject matter of claim 26 is not suggested by the Luczycki patent, even when considered in combination with the Shur patent.


Reconsideration and withdrawal of the rejection, and allowance of all pending claims, are respectfully requested.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

Date: September 24, 2003

By: _____



James A. LaBarre

Registration No. 28,632

P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620